

AUTOMATIC PROVISIONING OF SERVICES BASED ON A HIGH LEVEL

DESCRIPTION AND AN INFRASTRUCTURE DESCRIPTION

CROSS REFERENCES

The present application is cross-referenced to application entitled ".Automatic Provisioning of Services Based on Declarative Descriptions Of The Resource Structure Of The Service," ", having docket number YOR920040031US1, even dated herewith, and which is included herein by reference in entirety for all purposes.

FIELD OF THE INVENTION

The present invention is directed to provisioning and managing computing services in a computing utility system, based on high level description of the characteristics and structure of the desired computing services and a representation of the computing utility infrastructure used as a platform to implement the aforementioned services.

BACKGROUND

The cost and complexity of managing IT infrastructure continues to grow rapidly. Several factors contribute to this trend. First, IT infrastructures today are based on a distributed network of heterogeneous platforms and applications. In such an environment, resources, their capabilities, and behavior, are represented differently. They are therefore harder to compare and reason about. Interdependencies between resources, in which one resource's behavior is affected by another one, are not well represented or understood. Administrative personnel need to exercise knowledge of every platform, application and network appliance used, as well of how they can be assembled together for a particular purpose.

1 Second, in a highly competitive business environment, businesses must respond quickly to
2 market changes. Such changes may impose new requirements on the IT infrastructure, such as
3 supporting new computing services or applications, upgrading resources, incorporating new ones,
4 or changing the network structure. Realizing these changes is often a manual, tedious, and error
5 prone process. In particular, as configuration changes are made, new IT management processes
6 need to be defined, and existing processes may need to be updated.

7 Finally, service providers are moving towards an SLA-based service delivery model in which the
8 set of resources allocated to a customer is dynamically adjusted based on workload and
9 performance. Reconfiguring infrastructure resources dynamically in response to customer needs
10 demands prompt attention from administrative personnel increasing operational cost. Therefore,
11 a clear requirement of businesses today is to reduce the cost of maintaining an IT infrastructure
12 by reducing the overall complexity and the level of required human operation.

13 A common approach to addressing these challenges is incorporating automation into the
14 operation of the system. Common tasks such as adding a server to a computing service when the
15 load increases, or installing software on a server, are automated, thereby reducing human
16 involvement, the time to complete the task, and the probability for human errors. Workflows are
17 often used as a vehicle for automation because they are particularly well suited for coordinating
18 the execution of a set of activities that are long lived, tracking progress of activities, and
19 incorporating human interaction where necessary. *Provisioning engines*, including a workflow
20 engine, and some useful set of workflows organized in some structured way, are emerging as a
21 means to achieve the goal of reducing the cost through automation.

22 Automating the operation of the infrastructure, even by utilizing a provisioning engine, does not
23 fully address the aforementioned challenges. Automation procedures are often specific to a
24 particular infrastructure, computing environment, and service. When coding an automation
25 procedure it is impossible to predict all future changes in the service or infrastructure. Therefore
26 statically defined automation procedures are likely to require change. For example, changing

1 from a one-tier to a two-tier architecture, or adding resources with new capabilities, can require a
2 complex re-implementation of the automation procedures. This task is further complicated by
3 the many interdependencies between resources that are often implicit, by the combinatorial large
4 number of possible allocations and configurations of a given set of resources, by the variety of
5 possible computing services with different requirements, and by the many and rapidly evolving
6 types of hardware and software resources. Therefore there is a need to be able to describe a
7 computing service independent of a particular infrastructure, to describe the resources in a
8 service provider's infrastructure and their interdependencies, and to automatically generate the
9 instructions to provision and manage the service on the resources in the infrastructure.

10 **SUMMARY OF THE INVENTION**

11 Thus, this invention provides a process for performing provisioning given a high level
12 description of the desired computing service characteristics and requirements, independent of the
13 infrastructure, and a separate description of the infrastructure elements. A method is provided
14 which determines how to assemble the desired service environment from the building blocks
15 available in the infrastructure, or how to change the composition of resources allocated to an
16 existing service environment to meet new requirements. Automatic generation of provisioning
17 instruction facilitates consistent implementations, and reduces error. The instructions could be
18 embodied in a form such as a workflow that would serve as input to a provisioning engine.

19 The method is comprised of two distinct steps. In the first step, a Concrete Model is generated
20 from a Service Environment Model and an Infrastructure Model. The Service Environment
21 Model is a description of the characteristics of the desired service, independent of the
22 infrastructure. The Infrastructure Model encapsulates knowledge on elements of the
23 infrastructure, including resource instances, resource types, resource configuration, capabilities
24 and constraints. The Concrete Model satisfies two important properties. First, it is implementable
25 over the infrastructure. In other words, it can be created using infrastructure elements. Second,
26 as a refinement of the Service Environment Model, an implementation of the Concrete Model
27 also satisfies the requirements and characteristics described in the Service Environment Model.

1 The second step of the method comprises processing a given Concrete Model to generate
2 provisioning instructions such that a resource structure that matches the description in the
3 Concrete Model can be created in the infrastructure by executing the provisioning instructions.

4 The method can be used to perform provisioning, including creating a new service environment,
5 destroying an existing one, modifying the combination of resources allocated to a computing
6 environment, modifying their configuration, or any combination of the above.

7 In an example embodiment, a method comprises generating a Concrete Model. The Concrete
8 Model describes a structure of resources implementable over a computing utility infrastructure,
9 and satisfying a set of service requirements, said step of generating comprising the steps of:
10 obtaining a Service Environment Model of a service environment, said Service Environment
11 Model describing a new desired state of said service environment; getting an Infrastructure
12 Model describing both resources and an organization of the resources in the computing utility
13 infrastructure, said Infrastructure Model is encapsulated in a knowledge subsystem; and forming
14 the Concrete Model describing a resource structure such that said Concrete Model refines the
15 Service Environment Model and is mappable to said knowledge subsystem .

16 In an example embodiment, a method further comprising employing said Concrete Model to
17 generate provisioning actions, said provisioning actions, when executed, create a resource
18 structure that matches the description in the Concrete Model, said resource structure satisfies said
19 new desired state of said service environment.

20 In some embodiments, the method includes employing the Concrete Model to generate a resource
21 manager that manages a collection of composite resources.

22 **BRIEF DESCRIPTION OF THE DRAWINGS**

1 The foregoing and other objects, aspects and advantages will be better understood from
2 the following detailed description of a preferred embodiment of the invention with reference to
3 the drawings, in which:

4 **Fig. 1** The system in which the invention is used;

5 **Fig. 2** High level view of the management elements that are described in this invention;

6 **Fig. 3** High level view of process to generate provisioning instructions;

7 **Fig. 4** An example of a Service Environment Model;

8 **Fig. 5** An example of two possible service environment implementations;

9 **Fig. 6** Knowledge that is a part of the Infrastructure Model and the entities in the knowledge;
10 subsystem in which it is encapsulated.;

11 **Fig. 7** An example of a Concrete Model;

12 **Fig. 8** An example of a system structure with a scientific computing cluster service environment;

13 **Fig. 9** An example of an execution of Dynamic Provisioning Engine which shows the Front End
14 Generation and Back End Generation in details;

15 **Fig. 10 and Fig. 10A** Merging of a Service Environment Model and a 1-tier best practices
16 catalog template

17 **Fig. 11** The Front End Generation process; and

18 **Fig. 12** The Back End Generation process.

19 \

20 **DEFINITION OF TERMS**

21 A **service environment (SE)** supports a computing service offered to a single customer. It
22 includes hardware and software resources such as servers, operating systems, and middleware
23 configured to provide the required service

24 A **service provider** manages multiple service environments. A service provider employs a set
25 of **resources** that may be used to provision and maintain service environments. A service
26 provider is expected to be a prime user of the present invention.

1 A **computing utility** is the system that is used in order to provide computing services to
2 customer, including infrastructure resources and management software. The present invention is
3 a component of a computing utility.

4 A **computing utility infrastructure** is the set of resources that are used by a computing utility
5 system, including resources allocated to service environments (e.g., servers), and supporting
6 resources (e.g., network elements).

7 A **Service Environment Model** is a description, using a formal language, e.g. XML, of a
8 desired structure and state of a set of resources. This description is a high level description
9 which may be implemented in multiple ways in different infrastructures, namely, it is
10 infrastructure independent.

11 A **Concrete Model** is a description, using a formal language, e.g., XML, of a resource structure.
12 It includes description of a set of resources, including constraints on values of their attributes,
13 and a set of relationships between these resources.

14 A **knowledge subsystem** is a set of objects that are used to represent resource instances and
15 relationships, configure resources and relationships (including life cycle operations), query their
16 state, and query their configuration capabilities. The knowledge subsystem encapsulates
17 knowledge of the current state of the computing utility infrastructure, and constraints and
18 capabilities including policy based best practices defined by a service provider.

19 An **Infrastructure Model** is the knowledge encapsulated by the knowledge subsystem.

20 A **base** resource is a resource that is atomic; it cannot be broken down into other resources. Any
21 resource that a service provider does not want to subdivide can be treated as a base resource.
22 Adding, removing, or modifying a base resource changes the overall capacity of the system. We
23 assume that both the number and types of the base resources change over time. The notion of
24 base resource is service provider specific. Resources may be physical resources such as servers
25 and switch ports, logical resources such as server groups, IP addresses, and software licenses, or
26 virtual resources such as virtual servers or virtual local area networks (VLANs).

1 A **composite** resource is built out of other resources. A composite resource may have an
2 associated set of **relationships** between its constituent resources. An example of a composite
3 resource is a Web site. It may be composed of some number of front-end servers, back end
4 servers, a load balancer, a set of IP addresses for the servers, Web server software, database
5 software, and licenses associated with the software. The set of base resources used to implement
6 the function of a composite resource can change over time, though not all composite resources
7 may have this capability.

8 **Provisioning** refers to any task of creating, allocating, removing or configuring base or
9 composite resources in an existing or a new service environment. Provisioning actions may
10 involve assembling base resources into composites, configuring network devices, installing
11 operating systems, application software, monitors, user accounts, and so on. Since a service
12 environment is modeled as a composite resource, provisioning also refers to the act of setting up
13 a new service environment, modifying it, or destroying it.

14 **DESCRIPTION OF THE INVENTION**

15 The invention provides a method for generating provisioning actions given a description of the
16 desired computing service characteristics and a separate description of the infrastructure
17 elements. It determines how to assemble the desired service environment from the building
18 blocks available in the infrastructure, or how to change the composition of resources allocated to
19 an existing service environment to meet new requirements. The provisioning actions that are
20 generated can be embodied in a form such as a workflow that would serve as input to a
21 provisioning engine.

22 The environment in which the present invention operates may possess any of the characteristics
23 listed below. Although the present invention is not dependent on these characteristics, the
24 method is general enough to handle such conditions and characteristics.

1 First, resources may be allocated to customers in combinations which are heterogeneous, may be
2 interdependent, and vary over time.

3 Second, the service environments provided to each customer may be different. For example, one
4 customer may be provided resources for a web site, and another for a scientific computing
5 cluster. Resource types, quantities, dependencies, and allocation patterns will thus vary between
6 customers.

7 Third, there can be multiple ways to construct a service environment from the resources in a
8 service provider's infrastructure. A customer may have preferences or requirements for
9 particular variations of a given service environment. A service provider may also have
10 operational constraints that dictate which variations are acceptable.

11 Fourth, the infrastructure varies between service providers. Further, for a given service provider,
12 the infrastructure varies over time. These variations can be a result of upgrades or additions to
13 the physical infrastructure.

14 The system in which the present invention is used is a distributed computer system which has
15 multiple computing resources interconnected via a network. A diagram of such a system, which
16 might be found in a hosting or data center, is shown in Figure 1. The computing resources in the
17 physical infrastructure include, but are not limited to, processors 101, storage 103, firewalls 105,
18 and software 107. The software can be operating systems, middleware or applications. In figure
19 1 the available software has been preconfigured into Business Processes, Business Applications,
20 or Service Environments 107. The hardware resources are connected by a network 109 as
21 indicated by the grid of lines interconnecting all of these resources. This network may be
22 configure into one or more tiers, where each tier is separated by a router or firewall. Software
23 resources are assigned to physical resources by a management infrastructure. In this environment
24 a subset of the resources are assigned to the management infrastructure. In Figure 1 these
25 resources 111, 113, and 115 are indicated by a dotted line surrounding them. The resources
26 assigned to the management infrastructure run the management software described in this

1 invention. This software manages the rest of the resources. The resources used for the
2 management infrastructure are not assigned to customers of the hosted environment. The
3 remaining resources are assigned to customers as necessary. It is expected that customers 117
4 will primarily receive service by connecting to their resources through the Internet 119.
5 However, customers can receive service if they are connected to their resources by any means,
6 such as through a direct connection to a managed resource.

7 Figure 2 is a high level view of the elements that are described in this invention. They can be
8 divided into three groups. The physical infrastructure 201 is the actual set of interrelated
9 resources as described in Figure 1. The knowledge subsystem 237 contains a set of management
10 entities that represent and control the resources. In particular, every resource is represented and
11 controlled by a service with some canonical interfaces, termed a *Resource Instance Service*
12 (RIS). RIS 203 in the figure represents resource 207 as indicated by the dashed arrow 205
13 connecting them. Relationships and dependencies between resources in the physical
14 infrastructure are represented, in the knowledge subsystem, by a canonical set of relationships
15 between the corresponding RISs. The relationship 211 between RISs 203 and 213 indicates an
16 interdependency between the corresponding resources 207 and 209. The nature of this
17 interdependency is indicated by data associated with the relationship, such as its type, not
18 shown in the figure. Composite resources can also be represented and controlled by a RIS. For
19 example, RIS 217 represents a composite resource which is a service environment (SE). A
20 relationship of type *federates* connects RIS 217 with the three RISs 203, 213, and 215, as the
21 arrows indicate. This relationship means that the three corresponding resources in the physical
22 infrastructure are allocated to the service environment that RIS 217 represents. *Resource*
23 *Managers (RMs)* represent and manage collections of resources of the same type. In particular,
24 they provide an operation that returns a handle to a free resource instance. RM 219 manages the
25 set of processing capacity resources. RMs 221, 223, and 241 manage the collections of storage
26 firewall, and software licenses resources, respectively. A composite resource type may also have
27 an RM that encapsulates the knowledge of how to build the composite resource from other
28 resources. An RM for a composite resource may use the RMs of the resources that comprise the
29 composite resource. For example, in the figure, RM 225 is an RM for a *secure storage*

1 composite resource; it uses RMs 221 and 223, as the arrows indicate. Together, RMs and RISs
2 encapsulate knowledge of resource capabilities, how resources can be changed, how they can be
3 associated with other resources, and what resources are currently free or allocated. In managing
4 resources, RMs and RISs implement methods that affect changes on the resources they manage.
5 These methods may be implemented in any way such as by scripts or by a provisioning engine.

6 The management subsystem 239 contains a *dynamic provisioning engine* (DPE) 227 which
7 receives requests in the form of a Service Environment Model that describe a desired state, or a
8 set of requirements on the state of a set of resources, or a service environment. The DPE
9 generates provisioning actions 233 for reaching a state that satisfies the requirements specified in
10 the Service Environment Model. Once these provisioning actions are executed, either by the
11 DPE, or by the DPE using a provisioning engine, they affect the state of the system. For
12 example, a request may describe requirements on a new service environment, or requirements on
13 an existing service environment. Processing of a request results in provisioning actions that
14 change the combination of resources allocated to a service environment. To generate the
15 provisioning actions the DPE queries the knowledge subsystem, as indicated by arrow 235, to
16 understand the state of the system and how it can be changed. The DPE generates provisioning
17 actions that contain invocations of operations on knowledge subsystem entities, thus execution of
18 a sequence of provisioning actions affects the state of the system (resources in the physical
19 infrastructure), only through interaction with the knowledge subsystem (arrow 231), and not
20 directly. The generation and execution of provisioning actions may be interleaved; to serve a
21 request, a sequence of provisioning actions may be generated and executed before the next
22 sequence is generated ; sequences of provisioning actions may be regenerated if its execution
23 fails.

24 The focus of the invention is the process employed by the DPE in order to generate and execute
25 sequences of provisioning actions to create, destroy, or change the state of service environments,
26 or any combination of resources, given a high level description of the newly desired state. By a
27 *state of a service environment* we mean the combination of resources allocated to it, and their
28 configuration. Using the process, a sequence of provisioning actions to create, destroy or change

1 resource structure of a composite resource can be generated and executed automatically and
2 dynamically.

3 It is important to note that the same method can be applied inside a RM for a composite resource.
4 Thus, the method can be applied to automatically generate a RM for a composite resource type
5 based on its definition using a Concrete Model. The generated RM provides a set of methods to
6 create, destroy, or modify a composite resource based on a Concrete Model that describes its
7 desired structure. The RM can be then used as any other RM by a higher level DPE. Such a
8 strategy distributes the DPE method across infrastructure entities. It also enables creation and
9 re-use of provisioning components in different levels of granularity.

10 Figure 3 shows a high level view of the two stage process of this invention as it would be
11 employed in the DPE. The input to the process is (1) a *Service Environment (SE) Model*,
12 describing a set of requirements and, (2) a description 303, termed *Infrastructure Model*, of the
13 computing utility infrastructure. The Infrastructure Model represents knowledge on
14 infrastructure resources, their constraints and capabilities. This knowledge is encapsulated in the
15 knowledge subsystem. The DPE may need to query the knowledge subsystem in order to obtain
16 the parts of the Infrastructure Model that are necessary for its function.

17 The process includes of two main steps. In the first step, termed *Front End Generation*, the
18 Service Environment Model is combined with the Infrastructure Model to form a *Concrete*
19 *Service Model* 305. The Concrete Model declaratively describes the structure that need to be
20 created on the infrastructure in order to reach a state where all of the requirements expressed in
21 the Service Environment Model are satisfied. The second step, termed *Back End Generation*,
22 receives as an input the Concrete Model and generates and executes provisioning actions for its
23 construction 307.

24 **The Models**

25 The invention is based on the observation that services environments can be implemented in

different ways depending on the infrastructure at hand. For example, a *Web Site* service environment may be implemented using a one- or two-tiered network architecture, using firewalls or VLANs to implement security, and, Apache or Web Sphere as a Web server. The person who defines requirements and properties of the service environment that needs to be created may not have the infrastructure expertise, or even the knowledge of which infrastructure is going to be used. Therefore, a clear requirement is to separate the *Service Environment Model*, which describes requirements on service environments, from the *Infrastructure Model*, which describes an infrastructure, the available resources and how they are related, or can be related to each other. Hereafter, we describe the examples of the various models that can be used in this invention.

The Service Environment (SE) Model

The Service Environment (SE) Model (301, in figure 3) is a high level description of a set of requirements on a desired state of a service environment, independent of infrastructure. For example, it may describe a high throughput computing cluster. The description may include required components, properties, and behavior. It may describe a set of resources, properties, and relationships that must exist for the service environment to function. Figure 4 is an example of an SE Model for a scientific computing cluster service. The service includes resources of two types; a master node and a set of worker nodes. The worker nodes are used to perform scientific computations, while the master node controls the computation and in particular distributes the work between the workers and collects the results. Users should have access to the master node. Node 401 represents the service environment itself. It has a *federates* relationship with a single master node 405 and 3 worker nodes 403 as indicated by the multiplicity on the arrows. The interdependency between the master and workers is expressed by the *uses* relationship between the corresponding nodes. Node 407 describes a property of the master node; it must be accessible to external users.

Because the Service Environment Model is infrastructure independent, it is a partial specification. That is, it does not fully describe how the service environment will be

1 implemented on an infrastructure. There may be degrees of freedom in the description that allow
2 the service environment to be realized in different ways using different sets of resources or even
3 different infrastructures. Figure 5 shows an example of two possible implementations of the
4 scientific computing cluster service from Figure 4. In Figure 5a, the infrastructure is able to
5 support a two-tiered network organization including a front end VLAN 513, and a back end
6 VLAN 515. The master node 503 of the cluster serves as a front end to the worker nodes 505; it
7 is connected to both front end and back end VLANs. Only the master node is accessible to the
8 customer (501). This organization requires, for example, a master node with two network
9 adapters. Figure 5b shows an implementation of the cluster as a single tier. In it, the master
10 node 509 and the worker nodes 511 are connected to a single front end VLAN 517. In this
11 implementation, all the nodes are accessible to the customer 507. In this implementation only a
12 single network adapter is required in each node in the cluster.

13 In an infrastructure that supports both network architectures, the SE Model described in Figure 4
14 may be implemented either as a one- or two- tier structure. In this case, the DPE may choose one
15 of the architectures based on other criteria. For example, if at a certain point in time only servers
16 with one network adapter are available, then the one-tier architecture would be chosen.

17 **The Infrastructure Model**

18 An Infrastructure Model (303, in Figure 3) describes the resources and the organization of
19 resources in the service provider's infrastructure. It includes the resource types and capabilities
20 and information on physical connectivity, such as the number of network adapters and the
21 position of firewalls. Note that an Infrastructure Model need not exist as a single static entity. In
22 the preferred embodiment the Infrastructure Model is encapsulated and distributed among entities
23 of the knowledge subsystem, such as RMs and RISs. In order to obtain parts of the Infrastructure
24 Model necessary for its function the DPE may query entities of the knowledge subsystem.
25 Figure 6 describes knowledge seen as part of the Infrastructure Model and the way it might be
26 distributed among knowledge subsystem entities. In the knowledge subsystem (619) RISs 603
27 and 605 represent resource instances 607 and 609, respectively. Relationships between RISs

1 represent operational dependencies between the corresponding resources. For example, the
2 relationship 601 of type *uses* between RIS 603 representing server 607 and RIS 605 representing
3 shared file system (SFS) 609 indicates that the server 607 uses the shared file system 609. A RIS
4 can be queried for its set of relationships, as well as for values of configuration attributes of the
5 resource it represents. RMs manage collections of resources of the same type. An RM can be
6 queried to obtain information about the collection, and to obtain information on a type level. For
7 example, server RM 611 encapsulates a server type model 613 that can be queried for possible
8 relationships of a server resource to other resources. The set of possible server relationships will
9 include the aforementioned *uses* relationship between a server and a shared file system, as well as
10 other relationships that the server may have. The Infrastructure Model also includes information
11 on infrastructure constraints and capabilities. Infrastructure constraints can be expressed using
12 rules, assertions or other mechanisms. We also use attributes on relationships (on a type or
13 instance basis) to express constraints.

14 Following are some mechanisms that can be used to describe such infrastructure constraints. A
15 *fixed* attribute on a relationship expresses that the relationship between the two corresponding
16 resources cannot be changed by a provisioning action. For example, in a wire-and-forget
17 environment, where resources are wired exactly once to a set of switches, and the wiring cannot
18 be changed, a *connects* relationship between a Network Interface Card (NIC) and a Switch Port
19 (SP) will have the *fixed* attribute.

20 Another mechanism that can be used to express constraints and capabilities of the infrastructure,
21 as well as best practices that the service provider wishes to enforce is the *best practices catalog*
22 (615). It includes a set of patterns that formally represent best practices for constructing service
23 environments over the infrastructure. Patterns are defined once by an infrastructure expert and
24 are used by the DPE in the Front End Generation process to transform an Service Environment
25 Model to a Concrete Model. A pattern may be represented as a graph structure, where a node
26 may represent a resource type or a place holder that can be replaced by a sub-topology that
27 includes a set of nodes and relationships. Patterns may be annotated with attributes such as
28 SECURE to indicate the properties that the pattern supports. Patterns may also include mapping

rules from abstract to concrete resources. External mechanisms to generate such mappings may also be incorporated, for example we can use a third party mechanism that expresses dependencies that an application has.

In Figure 6, pattern 617 represents the structure of a 1-tier network architecture. In this pattern a service node represents any service (e.g., a scientific computation cluster service, or an e-commerce service). The server node represents any server (e.g., the master node or the worker node from Figure 4). The same idea can be used to define a 2-tier pattern, or any n -tier pattern. The DPE uses such patterns in the Front End Generation stage when the SE model is refined based on infrastructure knowledge to create a Concrete Model.

The Concrete Model

The *Concrete Model* (305 in Figure 3) is the output of the DPE's Front End Generation stage and the input to the Back End Generation stage. It is constructed by refining the Service Environment Model using the Infrastructure Model, therefore every resource structure that satisfies the requirements in the Concrete Model also satisfies the requirements expressed in the Service Environment Model. The goal of the Back End Generation stage is to generate such resource structure on the computing utility infrastructure. The refinement process uses the Infrastructure Model to ensure that the resulting Concrete Model is indeed implementable on the computing utility infrastructure.

The Concrete Model declaratively describes a structure of a set of resources that are associated with the service environment. More specifically, in a Concrete Model nodes represent resources, and requirements on the state of these resources. Edges represent relationships between resources. Every edge is associated with a list of attributes that describes the nature of the relationship. A node may contain a set of constraints on values of attributes of the resource that it represents. The values of some of these attributes are *fixed*, namely, they cannot be changed in a resource. Therefore, the constraints on these attributes are used as selection criteria for a resource that will serve the role of this node in the final resource structure that implements the

Concrete Model in the end of the provisioning process.

As in the Infrastructure Model, relationships between nodes in the Concrete Model may be *fixed* or *dynamic*; a *fixed* relationship cannot be changed; they reflect fixed infrastructure structures and operational constraints. Thus, such relationships must be taken into account in the selection of the resources. For example, if a Concrete Model contains a server node with a *fixed contains* relationship with three NICs then only a server with (at least) three NICs can be selected for this node. A *dynamic* relationship can be established by invoking a low level automation procedure on one (or more) of the knowledge subsystem entities. For example a dynamic *connects* relationship may be established between switch port and VLAN resources by programmatically configuring switches or routers.

In an advantageous embodiment, an edge, representing a relationship, is associated with a set of attributes that describe the nature of the relationship. Attributes describe the type of relationship (e.g., *federates*), and whether it is *fixed* or *dynamic*. A *color* attribute with value *green* denotes that the relationship must exist between the corresponding resources. The same attribute with value *red* denotes that the relationship must not exist between the corresponding resources.

While the Concrete Model shares many similarities with the Service Environment Model, there is one property that the Concrete Model should satisfy; it should be *mappable* onto the knowledge subsystem. More specifically, every node in the Concrete Model that represents a resource has to be mappable, either directly or indirectly, to either an RM or an RIS. Every relationship has to be mappable to an automation procedure to establish it (or un-establish it). The meaning of this property is that all of the high level concepts that were part of the Service Environment Model are now refined to a structure that is implementable using the knowledge subsystem. The most common case is that a node representing a resource is mapped to a RM for this type of resource. The mapping might be indirect; if one resource, say a server, has a *fixed contains* relationship with a different resource, say a NIC, then only the server node in the Concrete Model needs to be mappable to a server RM. In the Back End Generation stage, a server RIS can be obtained from the server RM, and a NIC RIS can then be obtained from the server RIS. Note that the NIC is

1 represented in the server type model encapsulated by the server RM. This condition will be
2 further explained when discussing the operation of the DPE. If only parts of the Concrete Model
3 can be mapped to the knowledge subsystem then the method of the invention can still be applied
4 to create a resource structure that matches in parts with the Concrete Model.

5 Figure 7 shows an example of a Concrete Model for a scientific computing cluster service that
6 corresponds to the system described in Figure 8. In this system every server (depicted in Figure 8
7 as a box) is connected to a designated Admin VLAN 801. A free server 813 has all of its other
8 NICs connected to a designated Free-pool VLAN 803. A service environment 813, is a scientific
9 computing cluster environment which includes a single master server 809, connected to both a
10 front end VLAN 805 and a back end VLAN 807, and a set of worker servers connected to the
11 back end VLAN. Both Master and Worker servers are also connected to the Admin VLAN.
12 This environment is a modification of that in Figure 5a that includes an administrative VLAN.

13 In Figure 7, the root node 701 is an object that represents the service environment itself; it
14 *federates* four resources: a master node 703, a set of zero or more worker nodes 707, a front end
15 VLAN 711, and a back end VLAN 709. The master node *contains* three network adapters
16 (NICs): one 715 connecting it through a switch port 727 to the Admin VLAN 729, one 705
17 connecting it through a switch port 723 to the front end VLAN 711, and one 713 connecting it
18 through a switch port 725 to the back end VLAN 709. The worker node *contains* two network
19 adapters: one 717 connecting it through a switch port 731 to the Admin VLAN, and one 719
20 connecting it through a switch port 721 to the back end VLAN. VLANs group switch ports;
21 Each VLAN is represented by a node which *contains* one or more switch ports.

22 The relationships between NICs and the servers and the NICs and the switch ports are *fixed*
23 relationships. They are defined when the example infrastructure is set up and this physical
24 connection is considered to be permanent. Such is also the case for the *contains* relationship
25 between the Admin NIC and a set of switch ports, as the servers in this example are to remain
26 permanently connected to the Admin VLAN. The dynamic relationships need to be established
27 by the DPE after the resources, represented by the nodes in the Concrete Model, are selected. An

example of a dynamic relationship is the relationship between the Back End VLAN 709 and the switch ports that it *contains* (725 and 721). This relationship is established by programmatically reconfiguring the switches. The DPE in the Back End Generation phase identifies such tasks and invokes the corresponding procedures in the knowledge subsystem to carry them out.

DPE Operation

Figure 9 is an example of an execution of the DPE which illustrates in more detail the two stage process for performing provisioning based on a high level model. The input to the Front End Generation stage is an Service Environment Model 901 and an Infrastructure Model 903. The Infrastructure Model represents information encapsulated in and obtained from the knowledge subsystem 925, as arrow 927 indicates. The Front End Generation is a sequence of refinement steps to the Service Environment Model based on the Infrastructure Model. Every refinement step produces an intermediary model 905 that serves as input to the next refinement step. In this example execution there are exactly two refinement steps 925 and 927. In the general cases there may be any number of refinement steps. The result of the Front End Generation is a Concrete Model 907 which serves as input to the Back End Generation stage. In the Back End Generation stage, resources corresponding to nodes in the Concrete Model are selected or created and relationships between them are established by interacting with the RMs in the knowledge subsystem. There may be multiple phases in which resources are selected or created and then configured to establish the corresponding relationships. In this example, after step 909 three resources are selected; the corresponding nodes are shown in black 917. In step 911 two relationships between the selected resources are established as depicted in 919. In step 913 the rest of the resources are selected as shown in 921. Finally, in step 915 all (dynamic) relationships between the resources are established as shown in 923. After the termination of the Back End Generation, a structure matching the Concrete Model structure is created in the knowledge subsystem. Here after we describe both stages in more details.

Front End Generation

1 In the Front End Generation stage, a Service Environment Model is refined based on an
2 Infrastructure Model to produce a Concrete Model. The Front End Generation is an iterative
3 process where every iteration further refines the intermediary model that is the result of the
4 previous iteration. In every iteration, a node in the model is selected and is replaced by a
5 subgraph structure which contains a set of nodes and edges. To refine a node entities in the
6 knowledge subsystem, and in particular the *best practices* catalog, are consulted. Several
7 structures in the *best practices* catalog are used. Following are some examples:

- 8
9 • The *best practices catalog* may contain a description of direct mapping from a node to a
10 sub-graph structure that refines it. For example, the best practices catalog may contain a
11 mapping from a single Web server node to a topology (sub-graph) that contains a server node
12 and an Apache software node, and a *hosts* relationship from the server node to the Apache
13 software node. Suppose that the DPE in the Front End Generation stage receives as an input
14 a Concrete Model that contains a Web server node. Once this node is selected, by consulting
15 with the *best practices* catalog, the DPE will replace the node with the aforementioned
16 sub-graph.
17
- 18 • The *best practices* catalog may contain patterns (a pattern is a representation of a family of
19 sub-graphs) that correspond to high level concepts such as SECURE, or, 1-tier, 2-tier, or,
20 *n*-tier. A node in an input model which describes a property, such as SECURE, once
21 selected, may be matched with a pattern annotated with a matching attribute. Consequently,
22 The Front End Generation will merge the pattern with the input model.

23 As an example, consider the Service Environment Model in Figure 4, and the 1-tier template
24 617 in Figure 6. Note that the keyword “external access” annotates the template 617.
25 Consequently, when node 407 in Figure 4 is selected in the Front End Generation stage, it
26 will be matched with pattern 617. The DPE will merge the two to produce the result
27 described in Figure 10A.

28 Figure 10 and Figure 10A describes the merging; for convenience, the Service Environment

Model and the *1-tier* pattern are shown again in Figure 10 (1031, and 1033, correspondingly). The result of the merging is shown in Figure 10A. First, the generic service node 1001 is replaced by the scientific compute cluster node 1003, which represents a specific service, resulting in node 1051 in the output model. Next, the generic server node 1005 in the pattern is replaced by two nodes, the Master node 1007 and the Worker node 1009, which both represent servers in specific roles, resulting in nodes 1059 and 1057 in the output model. The reasoning is that the generic server node in the pattern represents requirements on all the servers in the environment. Thus, both the Master server and the Worker server must satisfy these requirements. Note that in the pattern the multiplicity of a server is star which means “any number”. The Service Environment Model is more concrete in specifying the number of worker servers (three) and the number of master servers (one). Thus, three and one are used as the multiplicity of the corresponding nodes in the resulting graph 1053. The result satisfies and refines both models. Since there is a 1-1 *contains* relationship between a server and a NIC in the pattern, the NIC node 1011 in the pattern is transformed to two nodes 1063, and 1065 in the resulting output model, one for the Master server node and one for the Worker server node. Note that, as the Worker node 1017 actually represents three servers, Worker NIC 1065 represent three NICs, each contained in a different server. Using the same reasoning, the switch port node 1021 is transformed to a Worker switch port node 1067 and a Master Switch Port node 1061. In the 1-tier pattern all of the Switch ports, represented as a single node 1021, are contained in a single VLAN 1027. Thus, the VLAN node 1055 is added to the resulting topology and both the Master SP node and the Worker SP node are contained in it. The only node which disappeared from the Service Environment Model in this transition is the node 1029 (“External Access”) which represents a concept that is implemented by the 1-tier pattern.

- The *best practices* catalog may also include an external (generally a third party) refinement mechanism in order to dynamically generate sub graph patterns. Instead of mapping a node to a pattern, the node is mapped to a reference to an external procedure that can be a part of any external system. In the refinement process, when such node is selected, the reference is followed and the external procedure is executed. The result is a sub-topology that replaces

1 the selected node.

2 As described above, the refinement uses well known graph substitution and merging techniques
3 for each refinement step. The process halts when no more refinements can be made, or when no
4 more refinement is needed for the resulting model to be sufficiently mappable to the knowledge
5 plane. At this stage, it must be possible to map nodes to knowledge subsystem entities (RMs and
6 RISs).

7 We say that a Concrete Model *is mappable* to a knowledge subsystem if all of the nodes
8 representing resource roles are mappable directly or indirectly to a knowledge subsystem entity
9 (RM or RIS). The termination condition of the Front End Generation process is that no more
10 refinement on any node can be performed. At this state if resulting Concrete Model is mappable
11 to the knowledge subsystem then the process can continue to the Back End Generation, otherwise
12 the process fails since the service cannot be implemented on the infrastructure at hand.

13 An extension of this algorithm, clear to any one who is familiar with the art, continues with the
14 Back End Generation process even if the Concrete Model is only partially mappable to the
15 knowledge subsystem. In this case the end result of the entire DPE process may only partially
16 satisfy the requirements in the Service Environment Model.

17 Figure 11 is a diagram that describes the Front End Generation process. The process starts at
18 1101. It proceeds to step 1103 which is the first step in a single refinement iteration. In this step
19 the nodes are ordered, and the first node is selected and a Boolean variable *refine* is set to the
20 value *false*. Step 1105 determines if the selected node can be refined. If so, the process
21 continues to 1107 where the selected node is refined and a Boolean variable *refine* is set to *true*.
22 Otherwise, the process continues to step 1111. Step 1111 determines if the selected node was the
23 last node or not. If not, the next node is selected at step 1113. The process then returns to step
24 1105 to continue refinement. If the node was the last node, the process continues to step 1109.
25 At step 1109, the value of the variable *refine* is checked. The refine variable indicates if there
26 was a node that was refined in the course of a single iteration. If the iteration did not refine any
27 node then the refinement process exits to step 1115. In this case, no more refinement can be

1 performed on any node. Otherwise, a new iteration is started by returning to step 1103. When
2 the iterative refinement process stops, it is checked whether the Concrete Model is mappable to
3 the knowledge subsystem at step 1115. Accordingly, the process fails in 1117, or proceeds to the
4 Back End Generation in 1119.

5 **Back End Generation**

6 In the second stage, the DPE generates and executes provisioning actions to create a resource
7 structure that matches the Concrete Model and satisfies the requirements described in the Service
8 Environment Model. A provisioning action sequence contains two types of provisioning actions:
9 an action to select a resource, and an action to configure a resource or relationship. Usually a
10 sequence corresponding to a single request will contain multiple sub-sequences, termed *phases*,
11 in which resources are selected and then configured. The number of phases depends on the
12 complexity of the problem.

13 A single phase of the Back End Generation process is described in Figure 12. Essentially, in a
14 single phase, starting in step 1201, a subset of the resources are selected and mapped to nodes in
15 the Concrete Model (1203). This matching defines a set of provisioning actions of two types: a
16 provisioning action to configure a resource, and a provisioning action to establish or un-establish
17 a dynamic relationship. A provisioning action may have preconditions. A provisioning action
18 can be executed only if its preconditions are satisfied. In step 1209 a provisioning action whose
19 preconditions are satisfied is executed. An execution of a provisioning action may have side
20 effects in the physical infrastructure. In step 1211 these side effects are reflected back in the
21 knowledge subsystem by creating or destroying the corresponding relationships or changing the
22 values of attributes in the corresponding Resource Instance Services.

23 If all provisioning actions are executed successfully (1213) then the current phase of the Back
24 End Generation process terminates successfully (1223). If all nodes of the Concrete Model were
25 matched in this or previous phases then the Back End Provisioning process terminates
26 successfully and the entire structure described in the Concrete Model is now built in the physical
27 infrastructure and reflected in the knowledge subsystem. Otherwise, another phase is performed.

1 In some case, a provisioning action whose preconditions are not satisfied exists (1213). In this
2 case, an action to satisfy an unsatisfied precondition is executed (1221). As in the previous case,
3 side effects need to be reflected in the knowledge subsystem (1211). If actions exist whose
4 preconditions are not satisfied, and there do not exist any precondition that can be satisfied by
5 executing an action (1219 and 1225) then the process fails (1227).

6 We now describe some aspects of the Back End Generation in more detail. In the matching step
7 1203, selection of resources is based on two conditions: when a node is matched with a
8 corresponding Resource Instance Service, the values of fixed attributes as defined in the
9 Resource Instance Service must satisfy constraints on these attributes defined in the node. In
10 addition, all edges that represent fixed relationships with the node as an endpoint in the Concrete
11 Model must match the set of fixed relationships of the matching Resource Instance Service in
12 terms of type of relationship, direction of relationships, and matching endpoint nodes. More
13 specifically, if a node **A** in the concrete model is connected by an edge to a node **B** in the
14 concrete model and annotates with type x then there has to be a relationship of type x between the
15 Resource Instance Service **A'** that is matched with the node **A** and the Resource Instance Service
16 **B'** that is matched with the node **B**, moreover the direction of the relationship should be identical
17 to the direction of the aforementioned edge. The matching algorithm works by interacting with
18 the RMs. For every node, an operation is executed on the corresponding RM to find and obtain a
19 set of Resources Instance Services that are potential match for the node. The RMs may accept
20 some selection criteria (in the form of constraints over values of attributes) which are defined in
21 the node and passed as parameters to a *find* operation. These selection criteria only serve for the
22 initial filtering. Additional filtering must be done by the DPE so that the matching condition
23 defined above is satisfied. The matching is intricate since when selecting a resource
24 (represented by a Resource Instance Service) it is not enough to look only at its immediate fixed
25 relationships; a selection of a resource may dictate selection of a different resource (with whom
26 it has a relationship) so the latter one must also have the correct set of fixed relationships
27 recursively. To do the matching, the DPE employs well known graph matching techniques.
28 These techniques backtrack and try the next possibility whenever a matching possibility fails.

1 Once the matching is complete, a subset of the nodes is mapped to Resource Instance Services,
2 such that the set of fixed relationships and attributes in the Concrete Model matches the
3 corresponding relationships and attributes in the knowledge subsystem. This includes fixed
4 relationships between nodes that are both matched in this phase, or between nodes one of which
5 is matched in this phase and the other was matched in previous phases. In 1203, if a nontrivial
6 matching (i.e., matching of size greater than 0) exists, then the process proceeds to configuring
7 the resources, starting at 1207, otherwise the process fails in 1205. A process may fail due to
8 many reasons. For example, it may not be possible to map the pattern defined by the set of fixed
9 relationships to the infrastructure at hand. For example, if a node representing a server has 3
10 fixed *contains* relationships with nodes representing NICs and all servers in the infrastructure
11 have fewer than 3 NICs.

12 If a matching cannot be found the entire process fails (1205). Since some allocation and
13 configuration actions may already have been performed, a compensation action needs to take
14 place to restore the system state (1229). This is possibly done by recursively calling the DPE
15 with a new request such that the new desired state is the original state before the current DPE
16 process started. Alternatively, the algorithm can be easily generalized to find a non-optimal
17 solution in which a resource structure similar but not identical to the description in the Concrete
18 Model is found.

19 Once Resource Instance Services are selected and mapped to a subset of the nodes in the model,
20 they are configured to establish the set of dynamic relationships described in the model and to
21 change values of dynamic attributes to satisfy the constraints defined in the Concrete Model.
22 This is done by interacting with knowledge subsystem entities that encapsulate the logic to
23 configure the resources. Different knowledge subsystem entities may encapsulate automation
24 procedures to establish (or un-establish) different relationships. For example, a RM for a
25 composite resource may encapsulate the knowledge to establish all relationships between
26 resources in the composite. The invention does not make any assumption on the architectural
27 location of these automation procedures. It only assumes that such low level automation

1 procedures exist, and that there exists a mapping, accessible to the DPE, between a relationship
2 and the automation procedure to establish or un-establish it, and between an attribute and the
3 operation to set its value. An automation procedure will typically receive as parameters the
4 handles for resources involved and configure them to implement the semantics of the relationship
5 or attribute. The DPE is responsible for updating the corresponding RISs with the information
6 on the established relationship or value of attribute.

7 As explained above, a matching of a set of Concrete Model nodes and a set of Resource Instance
8 Services defines a set of provisioning actions that must be executed. The set includes two types
9 of provisioning actions: to configure a node and to establish or un-establish a relationship.
10 Specifically, for every node in the Concrete Model, for every attribute whose value is different
11 then the value of the attribute in the matching Resource Instance Service, a provisioning action
12 must be performed on the Resource Instance Service to change the value of the attribute. For
13 every edge representing a dynamic relationship in the Concrete Model, a provisioning action
14 must be performed to establish the relationship in the knowledge subsystem if the relationship do
15 not exist. For every red edge in the Concrete Model that represents a relationship that must not
16 exist, if a corresponding relationship exists in the knowledge subsystem it must be
17 un-established.

18 For example, a resource, say of type Web server, may have an attribute *state* whose value in the
19 Concrete Model is set to *started*. In the knowledge subsystem the values of the *state* attribute
20 may be *created*. An provisioning action must be performed to change the value from *created* to
21 *started*. The process can work as follows. An operation *setAttribute* is invoked on the Resource
22 Instance Service with parameters that are the name of attribute and the new value (“state”, and,
23 “started”, correspondingly). This operation triggers an automation procedure which affects the
24 physical infrastructure by starting the Web server that is represented by the Resource Instance
25 Service. For a relationships, consider as an example a relationship of type *contains* between a
26 Switch Port and a VLAN. An *establishRelationship* provisioning action can be invoked in the
27 knowledge subsystem that will trigger an automation procedure that affects the physical
28 infrastructure by programmatically configuring the switch to move the designated switch port in

1 the designated VLAN.

2 Back in Figure 12, the configuration is a 3 step process; first, in 1207 all dynamic relationships
3 and attributes between matched resources are collected by analyzing the Concrete Model and
4 mapped to the corresponding provisioning action, second, in 1209, a provisioning action whose
5 preconditions are satisfied is executed. As mentioned, a provisioning action may have side
6 effects that need to be reflected back in the knowledge subsystem. A side effect is any state
7 change that is beyond the property that is the target for which the automation procedure was
8 invoked. An example of a side effect of the operation to start a Webserver, described above, may
9 be the creation of a *use* relationship between the Webserver and a *database* resource. In this
10 case, step 1211 involves updating the knowledge subsystem with the aforementioned use
11 relationships. Side effects may be modeled and described in an inspectable way for every
12 operation in the knowledge subsystem that affects the physical infrastructure. The Front End
13 Generation may collect and include information about side effects in the Concrete Model.
14 Alternatively, they may be discovered by a different *discovery component* after the operation is
15 executed. In the later case once a provisioning action is executed, the discovery component is
16 executed and its output is used to update the knowledge subsystem in Step 1211. In some settings
17 the users of the system may decide that some relationships are not important for the management
18 of the system and they can be ignored all together.

19 Although the present invention may be employed by many types of entities, it is particularly
20 useful for use by a service provider, an enterprise owning an infrastructure used for running at
21 least one application, a customer of a service provider, a company owning an IT infrastructure,
22 and a utility provider.

23 Thus the present invention includes an apparatus comprising means for generating a Concrete
24 Model. The Concrete Model describes a structure of resources implementable over a computing
25 utility infrastructure, and satisfying a set of service requirements. The step of generating
26 comprising the steps of: means for obtaining a Service Environment Model of a service
27 environment, where the Service Environment Model describes a new desired state of the service

environment; means for getting an Infrastructure Model describing both resources and an organization of the resources in the computing utility infrastructure, the Infrastructure Model is encapsulated in a knowledge subsystem, and means for forming the Concrete Model describing a resource structure such that the Concrete Model refines the Service Environment Model and is mappable to the knowledge subsystem .

In some embodiments, the apparatus includes means for employing the Concrete Model to generate provisioning actions, the provisioning actions, when executed, create a resource structure that matches the description in the Concrete Model, the resource structure satisfies the new desired state of the service environment.

Variations described for the present invention can be realized in any combination desirable for each particular application. Thus particular limitations, and/or embodiment enhancements described herein, which may have particular advantages to a particular application need not be used for all applications. Also, not all limitations need be implemented in methods, systems and/or apparatus including one or more concepts of the present invention.

The present invention can be realized in hardware, software, or a combination of hardware and software. A visualization tool according to the present invention can be realized in a centralized fashion in one computer system, or in a distributed fashion where different elements are spread across several interconnected computer systems. Any kind of computer system - or other apparatus adapted for carrying out the methods and/or functions described herein - is suitable. A typical combination of hardware and software could be a general purpose computer system with a computer program that, when being loaded and executed, controls the computer system such that it carries out the methods described herein. The present invention can also be embedded in a computer program product, which comprises all the features enabling the implementation of the methods described herein, and which - when loaded in a computer system - is able to carry out these methods.

Computer program means or computer program in the present context include any expression, in

1 any language, code or notation, of a set of instructions intended to cause a system having an
2 information processing capability to perform a particular function either directly or after
3 conversion to another language, code or notation, and/or reproduction in a different material
4 form.

5 Thus the invention includes an article of manufacture which comprises a computer usable
6 medium having computer readable program code means embodied therein for causing a function
7 described above. The computer readable program code means in the article of manufacture
8 comprises computer readable program code means for causing a computer to effect the steps of a
9 method of this invention. Similarly, the present invention may be implemented as a computer
10 program product comprising a computer usable medium having computer readable program code
11 means embodied therein for causing a a function described above. The computer readable
12 program code means in the computer program product comprising computer readable program
13 code means for causing a computer to effect one or more functions of this invention.

14 Furthermore, the present invention may be implemented as a program storage device readable by
15 machine, tangibly embodying a program of instructions executable by the machine to perform
16 method steps for causing one or more functions of this invention.

17 It is noted that the foregoing has outlined some of the more pertinent objects and embodiments of
18 the present invention. This invention may be used for many applications. Thus, although the
19 description is made for particular arrangements and methods, the intent and concept of the
20 invention is suitable and applicable to other arrangements and applications. It will be clear to
21 those skilled in the art that modifications to the disclosed embodiments can be effected without
22 departing from the spirit and scope of the invention. The described embodiments ought to be
23 construed to be merely illustrative of some of the more prominent features and applications of the
24 invention. Other beneficial results can be realized by applying the disclosed invention in a
25 different manner or modifying the invention in ways known to those familiar with the art.